

CS-BWAMEM: A fast and scalable read aligner at the cloud scale for whole genome sequencing

Yu-Ting Chen^{1,*}, Jason Cong¹, Jie Lei¹, Sen Li¹, Myron Peto², Paul Spellman², Peng Wei¹, and Peipei Zhou¹

¹Computer Science Department, University of California, Los Angeles, CA 90095, USA

²Department of Molecular and Medical Genetics, Oregon Health and Science University, Portland, OR 97239, USA

ABSTRACT

Summary: The deep-coverage whole-genome sequencing (WGS) can generate billions of reads to be sequenced. It is time consuming for state-of-the-art aligners, such as BWA-MEM, to align the tremendous number of reads onto the reference genome. Inherently, the reads can be aligned using a massively parallel approach, and the alignment process should not be bounded by the limited number of computing cores of a single server. We present cloud-scale BWAMEM (CS-BWAMEM), an ultrafast and highly scalable aligner built on top of cloud infrastructures. It leverages the abundant computing resources in a public or private cloud to fully exploit the parallelism obtained from the enormous number of reads. With CS-BWAMEM, the pair-end whole-genome reads (30x) can be aligned within 80 minutes in a 25-node cluster with 300 cores.

Contact: ytchen@cs.ucla.edu

1 INTRODUCTION

Aligning reads to the reference genome is the first and one of the most time-consuming steps in many genomics pipelines, such as the pipeline for variant calling (McKenna *et al.*, 2010). The state-of-the-art aligners, such as BWA-SW (Li and Durbin, 2010), Bowtie2 (Langmead and Salzberg, 2012), and BWA-MEM (Li, 2013), can perform gapped-read alignment while considering pairing for pair-end reads. Today's next-generation sequencers can generate billions of reads in one run. Inherently, the alignment of a read is independent from the other reads. For the deep-coverage WGS data, the alignment takes more than 10 hours if we use BWA-MEM, which can only run on a single server. Since the computing resources are limited in a single server, we cannot fully exploit the huge parallelism arising from the billions of reads. Therefore, we develop a new aligner, CS-BWAMEM, on top of the cloud infrastructures and leverage the abundant computing resources in a computer cluster to accelerate the alignment process.

CS-BWAMEM has the following advantages compared to the state-of-the-art aligners. First, CS-BWAMEM is fast – 30x WGS data (300GB FASTQ reads) can be aligned within 80 minutes in a 25-node cluster. Second, CS-BWAMEM can provide scalable speedup. Users can decide the number of nodes in the cluster based on their performance demand. Third, CS-BWAMEM is validated with BWA-MEM and can obtain almost identical alignment quality.

2 METHODS

MapReduce is one of the most frequently used programming models for processing distributed data in a cloud. Apache Spark (Zaharia *et al.*, 2012) implements a novel and efficient MapReduce system that caches intermediate data, which can be reused in future steps, in memory to avoid unnecessary slow disk I/O accesses. The data in a cloud is usually stored in a distributed file system, such as HDFS (Shvachko *et al.*, 2010). In CS-BWAMEM, Spark is used as our computing framework, while HDFS is used to store WGS data.

CS-BWAMEM provides two major functions. First, the large FASTQ files, stored in the local Linux file system, are partitioned automatically into small file fragments and uploaded to HDFS. Second, the small FASTQ fragments are fetched from HDFS and processed in parallel on the allocated nodes in the cloud. For the CS-BWAMEM design, we reimplement the BWA-MEM algorithm using the MapReduce programming model. Like the two-stage design of BWA-MEM, CS-BWAMEM performs alignment in the first MapReduce stage while conducting pairing and generating aligned SAM files in the second stage. The intermediate data between the two stages are cached in memory to reduce the runtime.

We use a cluster with 25 nodes for evaluation. Each node equips with two 6-core, hyper-threading CPUs and 64GB RAM¹. CentOS 6.6, Apache Spark 1.3.1 and Apache Hadoop 2.5.2 (HDFS) are deployed. Figure 1(a) shows the runtime of aligning 30x WGS data (about 300GB). CS-BWAMEM scales well when the number of nodes increases. With 25 nodes, CS-BWAMEM can complete alignment within 80 minutes. For the alignment quality validation, we compare the final alignment of the data cleanup step suggested in GATK best practices for DNaseq for both aligners: CS-BWAMEM and BWA-MEM. Both aligners can successfully align 99.59% reads from a buffy exome sample. Figure 1(b) shows that CS-BWAMEM has almost identical numbers of aligned reads to BWA-MEM for all mapping quality scores, but provides a 7x speedup over BWA-MEM running on a single multi-core Intel Xeon server (Figure 1(a)).




Fig. 1. Results comparison between CS-BWAMEM and BWA-MEM on (a) runtime, and (b) the number of aligned reads according to the mapping quality scores on a buffy sample.

REFERENCES

- Langmead, B. and Salzberg, S. (2012). Fast gapped-read alignment with bowtie 2. *Nature Method*, **9**, 357–359.
- Li, H. (2013). Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. *ArXiv e-prints*.
- Li, H. and Durbin, R. (2010). Fast and accurate long-read alignment with Burrows-Wheeler transform. *Bioinformatics*, **26**(5), 589–595.
- McKenna, A., Hanna, M., Banks, E., Sivachenko, A., Cibulskis, K., Kernysky, A., Garimella, K., Altshuler, D., Gabriel, S., Daly, M., and DePristo, M. A. (2010). The genome analysis toolkit: a mapreduce framework for analyzing next-generation dna sequencing data. *Genome Res*, **20**(9), 1297–303.
- Shvachko, K., Kuang, H., Radia, S., and Chansler, R. (2010). The hadoop distributed file system. In *MSST, 2010*, pages 1–10.
- Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauley, M., Franklin, M. J., Shenker, S., and Stoica, I. (2012). Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *NSDI, 2012*, pages 2–2.

*to whom correspondence should be addressed

¹ We have CPUs with different number of cores. The aggregated computing power of our cluster is approximately 25 nodes.